

Trees of nuclei and bounds on the number of triangulations of the 3-ball

P. Collet¹, J.-P. Eckmann^{2,3} and M. Younan²

¹Centre de Physique Théorique, CNRS UMR 7644,
Ecole Polytechnique, F-91128 Palaiseau Cedex (France)

²Département de Physique Théorique, Université de Genève,
CH-1211 Genève 4 (Switzerland)

³Section de Mathématiques, Université de Genève, CH-1211
Genève 4 (Switzerland)

Abstract

Based on the work of Durhuus-Jónsson and Benedetti-Ziegler, we revisit the question of the number of triangulations of the 3-ball. We introduce a notion of nucleus (a triangulation of the 3-ball without internal nodes, and with each internal face having at most 1 external edge). We show that every triangulation can be built from trees of nuclei. This leads to a new reformulation of Gromov's question: We show that if the number of rooted nuclei with t tetrahedra has a bound of the form C^t , then the number of rooted triangulations with t tetrahedra is bounded by C_*^t .

Contents

1	Introduction	2
1.1	The method	3
1.2	Comparison with 2d	4
2	General definitions and notations	5
2.1	Internal and external objects, flowers	5
2.2	Notation and flowers	6
3	Some geometrical considerations: Two-colored paths in a triangulation	7
4	Part I: Reducing any triangulation into a set of nuclei	9
4.1	The elementary moves	9
4.2	Summary	12
4.3	Removing internal nodes	12
4.4	Reducing a triangulation with no internal nodes into a set of nuclei	21
5	Part II: Bounding the number of triangulations	21
5.1	Rooted triangulations	21
5.2	Trees of nuclei	22

5.3	Bound on triangulations	25
5.4	Combining the bounds	28

References	31
-------------------	-----------

1 Introduction

In this paper, we study the question of the number of triangulations of the 3-ball by tetrahedra. The case of the 2-ball was exactly solved by Tutte in [16]. He showed in particular that the number of rooted triangulations of the 2-sphere with N vertices is $\mathcal{O}(1) N^{-5/2} (256/27)^N$. It is natural to ask if analogous bounds are true in higher dimension. Such results could have applications in models of Statistical Mechanics (foams [14], quantum gravity [2], or glassy dynamics [1, 3, 8, 9]) where the exponential rate of growth can be interpreted as an entropy. In [12], Gromov asked whether the number of triangulations of the 3-sphere is bounded by C^N for some constant C when there are N tetrahedra (facets) in the triangulation. To date, this question remains open. However Pfeifle and Ziegler proved in [15] a super exponential lower bound for the number of triangulations of the 3-ball as a function of the number of vertices. This does not answer negatively Gromov’s question (which is in terms of the number of tetrahedra) but makes the problem of proving an exponential bound in terms of the number of tetrahedra even more challenging.

There are several studies in the direction of answering the question, which we summarize now. In [7], Durhuus and Jónsson gave the construction of a class of triangulations for which they could show a bound of the form C^N . These triangulations are obtained by building a tree of tetrahedra, which is obtained by starting from a root tetrahedron and attaching tetrahedra to its faces, and then attaching further tetrahedra to the new open faces. Each tetrahedron is attached to the tree with just one face. It is a common feature of tree-like constructions that they lead to bounds of the form C^N : The prime example in our context is of course the celebrated work of Tutte [16] mentioned above. Coming back to Durhuus and Jónsson, once the tree is constructed, they now collapse adjacent faces of the tree in such a way that at the end of the procedure a triangulation of the 3-sphere is obtained. Their main result says that the number of ways in which to do this is again exponentially bounded. In this way, they construct a set of triangulations of the 3-sphere with tetrahedra which is exponentially bounded. They ask whether these are all possible triangulations.

In a later development, Benedetti and Ziegler [4], show that the Durhuus and Jónsson construction, which they call “locally constructible” (LC), does *not* capture all triangulations of the 3-sphere. Namely, they show that a 3-sphere with a 3-complicated knot (made by tetrahedra) is not LC. They also carefully discuss relations between LC and other classes of constructibility.

In the present paper, we define a larger class of triangulations, with a construction similar to that of Durhuus and Jónsson, but which uses more general basic elements than the simple tetrahedron, which we call *nuclei*. We prefer to work with 3-balls, and bounds on 3-spheres can be obtained from a bound on triangulations of a tetrahedron. This is usually done by removing a tetrahedron from the 3-sphere (see for example [4, Section 3]).

Nuclei are defined as triangulations of the 3-ball with the following special properties:

1. They have no internal nodes.
2. Internal faces have at most *one* external edge.

Obviously, the tetrahedron is a nucleus. The Furch-Bing ball [10], [5] and [13] and the Bing 2-room house [5] and [13], which are not nuclei, can be reduced by our procedure to one non-trivial nucleus, each. The smallest non-trivial nucleus we know of, given in Table 1, has 12 nodes, and 37 tetrahedra, of which 17 have no external face. Nodes are numbered from 1 to 12, and Table 1 gives a list of the 37 tetrahedra.

1	3	4	10	1	3	5	10	1	3	5	11	1	4	6	10	1	5	7	8
1	5	7	10	1	5	8	11	1	6	7	8	1	6	7	10	2	3	5	9
2	3	5	11	2	3	8	9	2	3	8	11	2	5	6	11	2	6	11	12
2	7	10	11	2	7	11	12	2	8	9	10	2	8	10	11	3	4	9	10
3	4	9	12	3	5	9	10	3	8	9	12	4	5	6	11	4	5	7	8
4	5	8	11	4	6	10	11	4	7	8	9	4	7	9	12	4	8	9	10
4	8	10	11	6	7	8	9	6	7	9	11	6	7	10	11	6	8	9	12
6	9	11	12	7	9	11	12												

Table 1: A nucleus with 12 nodes, and 37 tetrahedra, of which 17 have no external face. If a tetrahedron has an external face, its 3 nodes are shown in boldface.

Our approach is two-fold: Top-down, and bottom-up. In the top-down approach, we define a set of elementary moves which reduces an arbitrary triangulation of the 3-ball into a tree of nuclei, which are glued together by pairs of faces, each such face with 3 external edges. The tree can then be cut into a disjoint union of nuclei by cutting along these faces. The construction always transforms 3-balls to unions of 3-balls, and is thus implementable on a computer.

In the bottom-up approach, we start with any tree whose nodes are arbitrary nuclei, and we construct 3-balls from it by gluing adequate faces together. Not all possible gluings lead to 3-balls, but including also some inadequate gluings still leads to good bounds. Again, the procedure can be programmed on a computer.

Our main result is Theorem 5.17. It says that *if the number $\varrho(t, f_s)$ of face-rooted nuclei with t tetrahedra and f_s external faces has a bound of the form $\varrho(t, f_s) \leq C^t$ then the number of rooted triangulations of the 3-ball with t tetrahedra, f external faces and n internal nodes is bounded by C_*^{t+f+n} .*

In particular, since obviously, $f \leq 4t$ and $n \leq 4t$, we would get a bound C_{**}^t .

In summary, our work bounds the number of triangulations in terms of the number of nuclei. Thus, we remain with a new, but hopefully simpler, open question about the problem posed by Gromov, namely does the number of face-rooted nuclei with t tetrahedra have an exponential bound in t ? While we do not have any mathematical statements about this problem, the methodology of the proof of Theorem 5.17 allows for quite extensive numerical experimentation. The most important insight from this experimentation is as follows: It seems that if T is a nucleus with a k -complicated knot (or even braid), then the addition of (at most) k cones and decomposition with our algorithm leads to a tree of *tetrahedra*. Note that the trefoil knot is 1-complicated. Furthermore, Goodrick [11] showed that the connected sum of k trefoil knots is at least k -complicated.

We have analyzed a certain number of classical examples, with the following findings summarized in Table 2.

1.1 The method

The bounds on the number of triangulations are obtained by studying a set of elementary moves, detailed in Sect. 4.1. These moves either decompose the triangulation in two disjoint pieces (by

Example	knot complication	# of cones added	ref.
Bing 2 room	no knot	1 cone	[5]
1 trefoil	1-complicated	1 cone	[10]
2 trefoils	2-complicated	1 cones	[4, Figure 3]
3 trefoils	3-complicated	2 cones	
4 trefoils	4-complicated	3 cones	
5 trefoils	5-complicated	3 cones	
figure 8	1-complicated	1 cones	
cinquefoil knot	1-complicated	1 cones	

Table 2: Experimental upper bound on the number of cones needed to decompose a triangulation into tetrahedra (For the definition of m -complicated, see [4]).

cutting along an interior face with 3 edges on the boundary, or taking away a tetrahedron with an external face and one internal node). Clearly, this leaves again two 3-balls on which we continue the decomposition. The other operations are “open” a ball along a carefully chosen edge (which we call “split-a-node-along-a-path”) or opening one face with 2 external edges. These operations *increase* the number of tetrahedra in the triangulation, but they prepare the moves in which the 3-ball can be cut, and the internal nodes can be eliminated. One of the main novelties of this construction is the observation that this can be done with *few* additional tetrahedra: This follows from a careful analysis of cuts of the 2-dimensional hemisphere attached to any external node. Since this is an important bound, we devote Sect. 3 to its proof. In Sect. 2, we introduce the (standard) terminology for the pieces of any triangulation. In Sect. 4 we combine the 4 moves described above to show how a general triangulation can be decomposed into a set of nuclei. In Sect. 5, we perform the bottom-up procedure and show how one bounds the number of triangulations of the 3-ball in terms of trees whose nodes are (rooted) nuclei, extending in this way the earlier work of [7] and [4].

1.2 Comparison with 2d

It is useful to compare our method to what can be done in 2d. In 2d we have a set of triangles. Any triangulation can be obtained in the following way: First, construct a tree of triangles, adding each triangle with only one face to the existing tree. This object has no internal nodes. Now, glue together adjacent faces of the tree, recursively. In this way one can obtain all triangulations of any polygon.

The inverse operation, while intuitively clear, is a little harder to describe, and we just sketch the procedure. Given any internal node x at distance 1 from the polygon, say connected to n we can split the edge (n, x) by doubling the node n into a pair n', n'' , so that the edges (n', x) and (x, n'') are now external edges and x is promoted to an external node. All internal nodes can recursively be brought to the surface in this way. We then have a tree, and the tree can be decomposed into triangles by cutting all internal edges with 2 external nodes. At the end, the basic objects are triangles.

Clearly, therefore, the basic objects in 2d are

- 2a) internal edges with 2 external nodes
- 2b) internal nodes (at distance 1) from the polygonal boundary

In 3d, there are many more possibilities, and our procedure will eliminate all those which can be eliminated. The ones which we can deal with are

- 3a) internal faces with 3 external edges: this corresponds to case 2a) above and will be cut by cut-a-3-face
- 3b) internal faces with 2 external edges, and therefore one internal edge with 2 external nodes. This resembles 2b) and is dealt with by open-a-2-face.
- 3c) an internal node x which is the tip of a tetrahedron t whose opposite face is external. One can just eliminate t and x becomes external. This is the second case which corresponds to 2b). We call this C0 later.
- 3d) an internal node x which is the corner of a face f whose opposite edge is external (but not C0). Again, a sub-case of 2b). This is dealt with split-a-node-along-a-path, and will be called C1.
- 3e) an internal node x which is the end of an edge e whose opposite end is external (but not C1). Again, a sub-case of 2b). This will be called C2 and reduced to C1 with split-a-node-along-a-path.

The elementary objects are those left over after all these decompositions are performed. In 2d, those objects are just triangles, which makes the counting possible. In 3d these are nuclei. Non-trivial nuclei exist, and they must carry the information about the complications of 3 dimensional topology, since all the other problems have been eliminated. In particular, internal faces of nuclei have 0 or 1 external edges.

2 General definitions and notations

2.1 Internal and external objects, flowers

To be precise, we redefine here some terminology which is common in the discussion of triangulations. We start with triangulations of S^2 . These will have f_s faces, n_s nodes and e_s edges, where the subscript s stands for “surface”. This triangulation is the boundary of a ball which is filled with tetrahedra, some of which have faces among the f_s external faces. We call this also a triangulation, and we say that t is the number of tetrahedra, f_{tot} the number of faces, e_{tot} the number of edges, and n_{tot} the number of nodes. Faces, edges, and nodes which are not among those of the triangulation of S^2 are called *internal*. It will be useful to observe that tetrahedra can have up to 4 external faces, internal faces can have up to 3 external edges, internal edges up to 2 external nodes. We will use the subscript i for internal objects.

Obviously,

$$f_{\text{tot}} = f_s + f_i, \quad e_{\text{tot}} = e_s + e_i, \quad n_{\text{tot}} = n_s + n_i.$$

From the Euler relations and trivial geometry, we have the relations

$$\begin{aligned} t - f_{\text{tot}} + e_{\text{tot}} - n_{\text{tot}} &= -1, \\ f_s - e_s + n_s &= 2, \\ 3f_s &= 2e_s, \\ 4t &= 2(f_{\text{tot}} - f_s) + f_s. \end{aligned} \tag{2.1}$$

This leaves us with 3 free variables, which we choose as

$$t, f_s, \text{ and } n_i.$$

Note that f_s is always even.

Definition 2.1. We use the term *f*-vector for the three variables $\langle t, f_s, n_i \rangle$ where $f_s \geq 4$.

2.2 Notation and flowers

We introduce some notation which we apply to triangulations and tetrahedrizations (which we also call triangulations when no confusion is possible):

- If n_1 and n_2 are 2 distinct nodes, then we denote by (n_1, n_2) the edge connecting the two.
- Similarly, if $n_i : i = 1, 2, 3$ are 3 distinct nodes, then (n_1, n_2, n_3) is the face (triangle) with these 3 corners.
- If e is an edge and n is a node not in e then (n, e) denotes the face (triangle) with the edge e and the node n .

This notation is easily generalized to the case where we consider simplices of dimension 3:

- If n is a node and f is a face not containing n , then (n, f) is the tetrahedron with f as a face and n as the opposite corner.
- Similarly, if e is an edge and $n_1, n_2 \notin e$ are 2 distinct nodes then (n_1, n_2, e) is the unique tetrahedron containing all of them.
- Finally, if e_1 and e_2 are two edges without common nodes, then (e_1, e_2) is the tetrahedron containing both edges.

Paths of nodes connected by edges will be denoted as $\gamma = [n_1, n_2, \dots, n_k]$ and the union of 2 disjoint paths γ_1, γ_2 (connected by one or both endpoints) will be denoted by $\gamma_1 \cup \gamma_2$.

We next define what we mean by *flowers*. Here, we adapt the conventions to the tetrahedrization of a triangulated sphere S^2 . Nodes, edges, and faces are called *external* if they lie entirely in S^2 . All others are called *internal*. Consider an external node n_* .¹ We define its 2 flowers:

- The *external flower* $\mathcal{E}(n_*)$ of n_* is the set of all edges e not containing n_* for which (n_*, e) is an external face. Clearly, $\mathcal{E}(n_*)$ is a polygon.
- The *internal hemisphere* $\mathcal{I}(n_*)$ of n_* is the set of all faces f not containing n_* for which (n_*, f) is a tetrahedron. It is easy to see that $\mathcal{I}(n_*)$ is a 2d triangulation whose boundary is the polygon $\mathcal{E}(n_*)$.

We will say that the external flower of an *internal* node x_* is empty. The internal (hemi-)sphere $\mathcal{I}(x_*)$ (or simply flower) of x_* is a triangulation of S^2 .

We also define the *external flower* $\mathcal{E}(e)$ of an *external edge* e as the 2 nodes n_1, n_2 for which (n_i, e) are 2 external faces. Similarly, the *internal hemisphere* $\mathcal{I}(e)$ of the *external edge* e is defined as the set of all edges e' such that (e, e') is a tetrahedron. By hypothesis, $\mathcal{I}(e)$ is a 1-d triangulation whose boundary is $\mathcal{E}(e)$. Note that there might be internal nodes at distance 1 from e which are not in $\mathcal{I}(e)$.

¹We use n_*, m_* and the like for external nodes, and x_*, y_*, \dots for internal ones.

3 Some geometrical considerations: Two-colored paths in a triangulation

We describe here properties of paths in a 2d triangulation of a polygon. These properties will play a crucial role when we will bound the effects of moving internal nodes of a 3d triangulation to the surface. However, they are totally independent of the remainder of the paper.

Lemma 3.1. *Let \mathcal{K} be a 2d triangulation of a p -gon P with n interior nodes. Then the number of interior edges in \mathcal{K} is $3n + p - 3$.*

Proof. The proof follows from the Euler relations and is left to the reader. \square

Lemma 3.2. *Consider a polygon P and let \mathcal{K} be any triangulation of P with $k > 0$ internal nodes. For each node $x \in \mathcal{K} \setminus P$, there are at least 3 simple disjoint paths in the interior of \mathcal{K} connecting it to 3 different points of P .*

Proof. Any triangulation of S^2 is 3-connected. Complete \mathcal{K} into a triangulation of S^2 by adding a cone over its boundary. Let m be the apex of the cone. Then there are at least 3 disjoint simple paths connecting x to m , [6]. Any such path must intersect P , and we take the first intersection point. \square

We assume now that the nodes of P are labeled.

Definition 3.3. *A triangulation \mathcal{K} is called admissible if the following conditions are met:*

K1: The boundary $\partial\mathcal{K}$ has at least 2 different labels.

K2: The nodes with a given label form one connected arc of $\partial\mathcal{K}$.

K3: The ends of any edge connecting 2 nodes of $\partial\mathcal{K}$ have different labels, unless the edge is in $\partial\mathcal{K}$.

The Fig. 1 illustrates the definition.




Figure 1: An illustration of the conditions K1)–K3). Left: since there is only one label, K1) is violated. Center: The region with label 1 is not connected; K2) is violated. Right: There is an internal link (red) connecting two nodes with the same label; K3) is violated.

We first need an auxiliary lemma. We will need the following information:




Figure 2: The 2 alternatives of finding a path connecting two different labels. Left: There is an interior path between a and b . Right: There is no such path, but then, one can always find an edge connecting two different labels (by K3), (not necessarily the same as a and b). The left panel also illustrates the necessity of choosing a shortest path. For example, choosing the magenta path, the dashed edge will violate K3) in the next step of the procedure.

Lemma 3.4. *Let \mathcal{K} be as above and let $P = \partial\mathcal{K}$. Given two boundary nodes a and b with different labels at least one of the two alternatives below holds:*

- 1) *There is a simple path γ joining a and b without any other node in P ,*
- 2) *There is an edge (x, y) joining the two pieces of $P \setminus \{a, b\}$.*

Postponing the proof of Lemma 3.4 we have

Proposition 3.5. *Assume \mathcal{K} is an admissible triangulation in the sense of Definition 3.3 with at least 2 triangles. Then, there exists a path γ along internal edges of \mathcal{K} which connects two points in $P = \partial\mathcal{K}$ with different labels. It cuts \mathcal{K} in two pieces \mathcal{K}_L and \mathcal{K}_R . The path γ can be chosen in such a way that labeling the new boundary piece (namely the interior nodes of γ) in \mathcal{K}_L and \mathcal{K}_R with a label different from the ones used so far, both \mathcal{K}_L and \mathcal{K}_R are again admissible.*

Proof. Let $P = \partial\mathcal{K}$. By admissibility, we know that not all nodes on P have the same label. Take nodes a and b with different labels and apply Lemma 3.4. If 2) holds, then we take γ as the edge connecting x and y . By K3) they have different labels. Otherwise, there is an interior path connecting a and b . We take a shortest path, γ .

Cutting along the path γ , we obtain 2 pieces \mathcal{K}_L and \mathcal{K}_R . If γ is just one edge then inspection shows that K1)–K3) hold. In the second case, K1) and K2) are obviously true by construction. Giving a new label, say L , to the interior nodes of γ , we have to show that there are no edges connecting any two non-consecutive nodes with label L . But if there were, the path would not be minimal. \square

Proof of Lemma 3.4. The reader may want to look at Fig. 2. Assume 1) does not hold. This means that one cannot draw 3 disjoint paths between a and b , as the middle one would satisfy 1). We can

take the two disjoint paths to go along the two boundary segments between a and b . By Menger's theorem [6] there must then be 2 nodes x and y (other than a or b) such that all paths from a to b must pass through at least one of them. Since the boundary paths are candidates, we see that x and y are in P , one per arc connecting a and b . Consider now the path from a to b along P which goes through x . Modify it so that instead of going through x it goes through the flower of x . We get a new path from a to b which does not go through x . This means that the new path goes through y implying that y is in the flower of x . Thus, x and y are connected by an edge.

This completes the proof. \square

4 Part I: Reducing any triangulation into a set of nuclei

4.1 The elementary moves

In this section we define the elementary moves which transform any triangulation into a (set of) nuclei. The first two moves, which we call *open-a-2-face* and *cut-a-3-face*, are used to transform any triangulation with no internal nodes into a set of nuclei, and the third and fourth move, which we call *remove-1-tetra* and *split-a-node-along-a-path*, are used to remove all internal nodes of a triangulation.

Henceforth, T will denote a triangulated 3-ball with f-vector $\langle t, f_s, n_i \rangle$.

4.1.1 Cut-a-3-face

Let (n_1, n_2, n_3) be an internal face with its 3 edges on the surface ∂T of T . Then, it cuts the 3-ball T into 2 distinct parts. We simply separate these 2 parts and we get 2 “smaller” 3-balls. In other words, we know that any triangulation is (uniquely) defined by the list of all its tetrahedra. We find the two lists corresponding to the tetrahedra which are on either side of the internal face in question, and we define 2 new 3-balls, each associated with one of these 2 lists. If $\langle t, f, n \rangle$, $\langle t_1, f_1, n_1 \rangle$ and $\langle t_2, f_2, n_2 \rangle$ are the f-vectors of the initial ball and the 2 new ones, then we have

$$t = t_1 + t_2, \quad f = f_1 + f_2 - 2, \quad n = n_1 + n_2.$$

4.1.2 Open-a-2-face

Consider 3 external nodes n_*, n_1, n_2 of T which form a triangle (n_*, n_1, n_2) . We assume that (n_*, n_1, n_2) is an internal face, with (n_1, n_2) an internal edge, and the two other edges external. Let \mathcal{I} and \mathcal{E} be the internal and external flower of the external node n_* . As we have already stated, \mathcal{I} is a triangulation of the polygon \mathcal{E} . By hypothesis, the edge (n_1, n_2) divides \mathcal{I} into 2 distinct sets of faces. The operation *open-a-2-face* consists in removing n_* and all tetrahedra attached to it, replacing it by $n_{*,1}$ and $n_{*,2}$ and attaching each of these 2 new nodes to all faces of one of the two parts of \mathcal{I} , see Fig. 3. *This operation transforms a triangulation of the 3-ball into a triangulation of the 3-ball.* If $\langle t, f_s, n_i \rangle$ and $\langle t', f'_s, n'_i \rangle$ are the f-vectors of the initial ball and the resulting ball, then we have

$$t' = t, \quad f'_s = f_s + 2, \quad n'_i = n_i.$$

We will say that the f-vector changes by $\langle 0, +2, 0 \rangle$.




Figure 3: Sketch of open-a-2-face.

4.1.3 Remove-1-tetra

Definition 4.1. A removable tetrahedron is any tetrahedron t with one internal node and one external face.

The operation remove-1-tetra is as follows: let $t_* = (x_*, n_1, n_2, n_3)$ be a removable tetrahedron with internal node x_* and external face (n_1, n_2, n_3) . We simply remove t_* and its external face; the internal node x_* , the 3 internal edges and the 3 internal faces of t_* all become external. The f-vector $\langle t, f_s, n_i \rangle$ changes to $\langle t - 1, f_s + 2, n_i - 1 \rangle$; the change of f-vector is $\langle -1, 2, -1 \rangle$.

4.1.4 Split-a-node-along-a-path, hemispheres and pieces

Consider an external node n_* of T and its internal hemisphere $\mathcal{I} = \mathcal{I}(n_*)$, see Fig. 4 for an illustration. By definition of a triangulation, \mathcal{I} is a 2d triangulation of a polygon.

Definition 4.2. A splitting path γ is any simple path which connects two different points on $\partial\mathcal{I}$ and contains no edge of $\partial\mathcal{I}$.

Let γ be a splitting path. Clearly it divides \mathcal{I} into 2 pieces \mathcal{K}_L and \mathcal{K}_R with $\mathcal{I} = \mathcal{K}_L \cup \mathcal{K}_R$ and $\mathcal{K}_L \cap \mathcal{K}_R = \gamma$.

The move split-a-node-along-a-path γ is defined as follows:

1. Remove the node n_* and all tetrahedra having n_* as a corner
2. Add 2 new nodes $n_{*,L}$ and $n_{*,R}$
3. For each face $f_* \in \mathcal{K}_L$ add the tetrahedron $(n_{*,L}, f_*)$
4. For each face $f_* \in \mathcal{K}_R$ add the tetrahedron $(n_{*,R}, f_*)$
5. For each edge $e \in \gamma$ add the tetrahedron $(n_{*,L}, n_{*,R}, e)$

Note that by construction, one of the nodes on $\partial\mathcal{K}_L$ is $n_{*,R}$, and the links in \mathcal{K}_L starting from $n_{*,R}$ reach (the image of) γ . Analogous statements hold for \mathcal{K}_R .

Definition 4.3. In the construction above, we refer to $\mathcal{K}(n_{*,L}) = \mathcal{K}_L$ as the left piece. It is simply the subgraph obtained from the hemisphere $\mathcal{I}(n_{*,L})$ after removing the cone connecting $n_{*,R}$ to every node of γ . Similarly, we define the right piece \mathcal{K}_R .

Remark 4.4. Hemispheres \mathcal{I} and pieces \mathcal{K} will play an important role in our construction. Some statements will be given for hemispheres, others for pieces and so it is important to be able to distinguish between the two definitions.

Remark 4.5. A splitting path γ is always associated with a hemisphere \mathcal{I} and not with a piece \mathcal{K} . We will see that, under some conditions, a simple path $\tilde{\gamma}$ connecting two nodes of the boundary of a piece \mathcal{K} can be extended into a splitting path γ .




Figure 4: The left panel shows the internal hemisphere $\mathcal{I}(n_*)$ of n_* . We split n_* into $n_{*,L}$ and $n_{*,R}$ along the green path γ . The other 2 panels show the internal hemispheres $\mathcal{I}(n_{*,L})$ and $\mathcal{I}(n_{*,R})$ of the 2 new nodes. Notice that each internal node of the green path γ is at distance 1 from $n_{*,R}$ resp. $n_{*,L}$. Also, the links leaving $n_{*,s}$, $s \in \{L, R\}$ have been added during the split.

Lemma 4.6. The move split-a-node-along-a-path transforms a 3-ball into a 3-ball. The f-vector $\langle t, f_s, n_i \rangle$ is mapped to $\langle t + |\gamma|, f_s + 2, n_i \rangle$, where $|\gamma|$ is the number of edges in γ .

The f-vector changes by $\langle |\gamma|, +2, 0 \rangle$. In particular, the number of tetrahedra *increases*. But we will show that this increase can be controlled.

Proof. The count of the f-vector is as follows: Removing and adding the tetrahedra in steps 1,3,4 above does not change their number. The number of external faces increases by two, namely the two external faces sharing the new edge $(n_{*,L}, n_{*,R})$. And each internal face (n_*, e) which connected n_* to an edge e in γ gives rise to a new tetrahedron $(n_{*,R}, n_{*,L}, e)$. There are $|\gamma|$ such faces and so the f-vector is seen to change by $\langle |\gamma|, +2, 0 \rangle$, as asserted. \square

4.2 Summary

In the sequel, we want to bound the effect of removing internal nodes, since our building blocks are the nuclei, which do not have any internal nodes. Eliminating the internal nodes will cost the addition of tetrahedra, and the issue here is how many are needed to obtain a ball without internal nodes. Internal nodes disappear when we perform the remove-1-tetra operation, and only then.

Before starting the bounds proper, we explain here the point of our construction, based on the evolution of the f-vectors $\langle t, f_s, n_i \rangle$. Open-a-2-face costs a change $\langle 0, 2, 0 \rangle$, and split-a-node-along-a-path costs $\langle |\gamma|, 2, 0 \rangle$, where $|\gamma|$ is the length of the path along which we cut. In principle, each path γ might have a length proportional to the number of nodes, which in turn would imply that the sum of the lengths of all paths exceeds $\mathcal{O}(n_{\text{tot}}^2)$. So one needs a strategy which improves this naive bound.

While we cut, new external edges appear, and also, new external edges appear when we remove a tetrahedron which costs $\langle -1, 2, -1 \rangle$. But it is only this operation which reduces the number of internal nodes. So, there are two opposing tendencies. One is the preparation of promoting an internal node into an external one, and it *adds* many tetrahedra, and the other is remove-1-tetra, which reduces the number of internal nodes by 1. The real issue is thus to bound the *number of added tetrahedra per removed internal node*. We will perform this bound in terms of the number e_s of internal edges. Our main result is Corollary 4.12 which says that the number of internal edges grows by no more than $C_\Delta(t + n_i)$. The Euler relations (2.1) allow to express t as a function of e_s , f_s , and n_i ,

$$t = e_i - n_i + f_s/2 - 1 .$$

Therefore, and since $n_i < 4t$ and $f_s < 4t$, Corollary 4.12 implies that the elimination of all n_i internal nodes leads to an f-vector of the form

$$\langle t, f_s, n_i \rangle \rightarrow \langle t', f'_s, 0 \rangle , \quad f'_s < C \cdot t , \quad t' < C \cdot t ,$$

with a finite constant C which is *independent* of the triangulation.

4.3 Removing internal nodes

4.3.1 Definitions and strategy

Given any triangulation, we define the *depth* D_x of a node x as the minimal number of connected edges needed to reach the boundary, starting from x . The strategy will consist in recursively reducing the depth of any internal node by 1. This is repeated until no internal nodes remain.

We classify an internal node x_* at depth 1 in 3 flavors, which we call C0–2:

C0: x_* is the internal node of a removable tetrahedron.

C1: x_* is not of type C0 but is in a face (x_*, n_*, m_*) where (n_*, m_*) is an external edge.

C2: x_* is neither of type C0 nor C1.

We enumerate the external nodes in an arbitrary order, leading to a list $\mathcal{L}_0 = \{n_{*,1}, \dots, n_{*,k}\}$. Similarly, for $d > 0$, we define \mathcal{L}_d as the nodes at depth d from the surface. Given an external node $n_* \in \mathcal{L}_0$, we consider its hemisphere $\mathcal{I}(n_*)$.

An internal node $x_* \in \mathcal{I}(n_*)$ of type C2 can (only) be *promoted* to an internal node of type C1 by drawing a path $\gamma \subset \mathcal{I}(n_*)$ that goes through it and splitting n_* into $n_{*,L}$ and $n_{*,R}$ along γ . Indeed, one easily sees that $n_{*,R} \in \mathcal{E}(n_{*,L})$ and that $(n_{*,L}, n_{*,R}, x_*)$ is a face.

In the same manner, we see that a node $x_* \in \mathcal{I}(n_*)$ of type C1 can be promoted into an internal node of type C0 by drawing a path $\gamma \subset \mathcal{I}(n_*)$ which contains the edge (x_*, y) . Here, y is the external node of the face (x_*, n_*, y) which defines x_* as a node of type C1. Splitting n_* along γ , the tetrahedron $(n_{*,L}, n_{*,R}, y, x_*)$ becomes removable.

Finally, any internal node of type C0 can be made external by simply removing one tetrahedron. The strategy is in 4 steps (3 sweeps). We set $\mathcal{L} = \mathcal{L}_0$.

- **Step 1 (Sweep C2→C1)** : We promote all the x_* of type C2 in the following order: For each $n_* \in \mathcal{L}$, we promote all internal nodes of $\mathcal{I}(n_*)$ of type C2 into internal nodes of type C1. We will show that this can be done in such a way that every internal edge of the triangulation $\mathcal{I}(n_*)$ belongs to at most 1 of the splitting paths (as defined in Sect. 4.1.4).
When this first step is complete, all internal nodes at depth 1 are of type C1 or C0. There appears a new set \mathcal{M} of external nodes containing the nodes of \mathcal{L} which were not split and new external nodes obtained by the splitting.
- **Step 2 (Sweep C1→C0)** : We promote all the x_* of type C1 in the following order: For each $n_* \in \mathcal{M}$, we promote all promotable internal nodes of $\mathcal{I}(n_*)$ of type C1 into internal nodes of type C0². We will show that this can be done in such a way that every internal edge of the triangulation $\mathcal{I}(n_*)$ belongs to at most 1 of the splitting paths (as defined in Sect. 4.1.4).
- **Step 3 (Sweep C0→external)** : Finally, we make each node of type C0 external by removing one tetrahedron.
- **Step 4** : At this point every internal node has been moved up one level of depth. In particular, we let \mathcal{L} denote those nodes which have moved to the surface in step 3. (If the current step is at level d , then this set equals \mathcal{L}_{d+1} .)
We continue until no internal nodes are left.

Since the depth of any node is bounded, the procedure will end after a finite number of recursive steps.

4.3.2 Reducing C2-nodes to C1-nodes

Given an external node $n_* \in \mathcal{L}$, we now describe in detail the recursive algorithm which promotes the internal nodes of type C2 in $\mathcal{I} = \mathcal{I}(n_*)$ to type C1. This is achieved by a succession of carefully chosen moves of type split-a-node-along-a-path.

Each of these cuts produces a “left” and a “right” piece, which are then cut again into left and right pieces, until only triangles remain. The pieces will be noted by $\mathcal{K}_{\mathbf{S}} = \mathcal{K}(n_{*,\mathbf{S}})$, where \mathbf{S} is a sequence of letters L and R which designate the successive choices of left and right.

Thus, we construct a binary tree of pieces (see Fig. 5). In detail:

²A node x_* of type C1 can be promoted to C0 only if it is connected to a node of $\mathcal{E}(n_*)$. This might not be true for all n_* for which $x_* \in \mathcal{I}(n_*)$ but there are at least two n_* for which x_* is promotable.

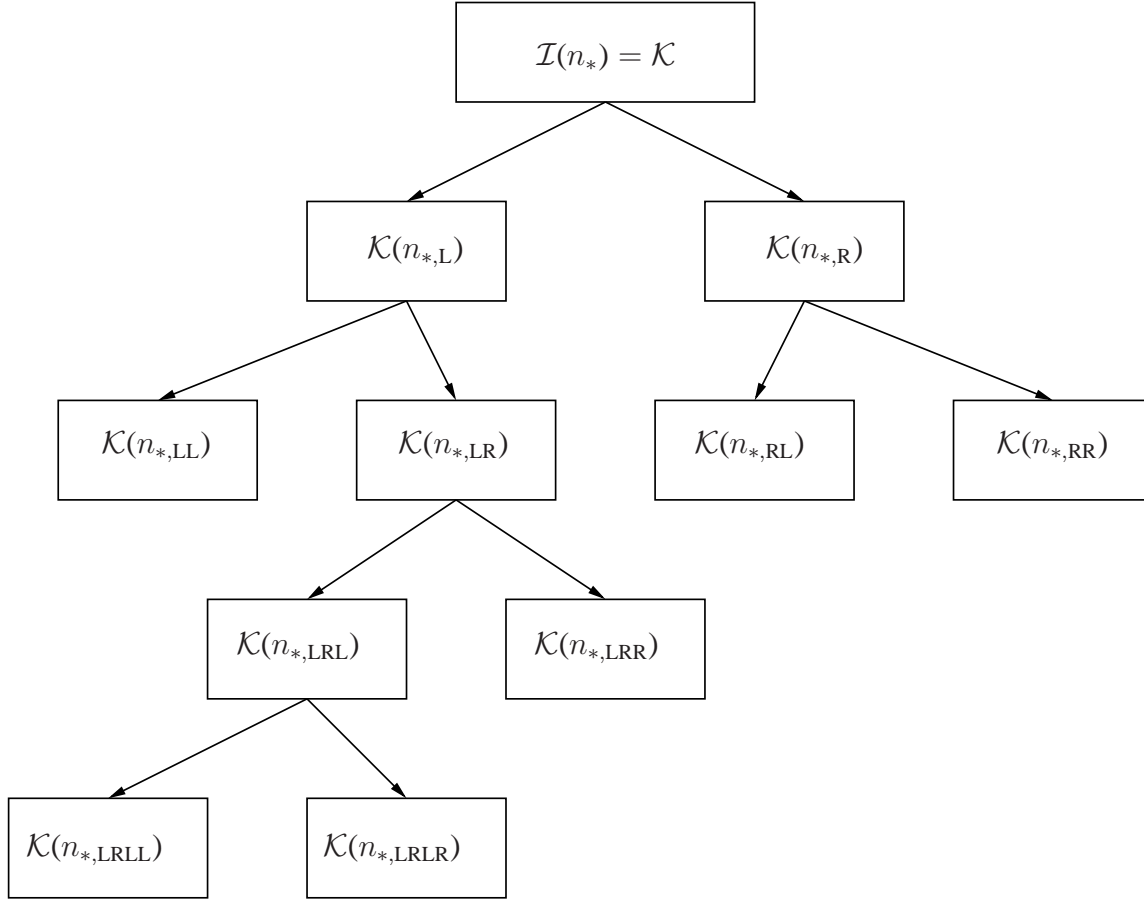


Figure 5: An example of a binary tree of pieces associated with the hemisphere of an external node n_* containing 6 triangles.

1. Label the nodes of $\partial\mathcal{I}$ from -1 to $-|\partial\mathcal{I}|$.
2. The hemisphere \mathcal{I} is an admissible triangulation in the sense of Definition 3.3. Proposition 3.5 implies the existence of a shortest path γ which connects two nodes of $\partial\mathcal{I}$ (with different labels). We choose this path γ .
3. After splitting along this path, \mathcal{I} is divided in two pieces, as shown in Fig. 4. The two pieces are called \mathcal{K}_L and \mathcal{K}_R . The splitting has replaced n_* by $n_{*,L}$ and $n_{*,R}$ and $\mathcal{I}(n_{*,L})$ is actually just \mathcal{K}_L with the cone between $n_{*,R}$ and γ added. This also means that $n_{*,R}$ is in the external flower $\mathcal{E}(n_{*,L})$ of $n_{*,L}$. Analogous terminology is used for the other half. At this point, \mathbf{S} is equal to L or R, and we continue with $\mathbf{S} = \mathbf{L}$ (and do later $\mathbf{S} = \mathbf{R}$).
4. If \mathcal{K}_S is a triangle, we are done (for this branch of the tree).
5. Label all nodes on $\partial\mathcal{K}_S$ which had no label with the label $\hat{\mathbf{S}}$, where $\hat{\mathbf{S}}$ is obtained from \mathbf{S} by exchanging the last letter, cf. Fig. 6. In this way, the newly labeled nodes are connected to $n_{*,\hat{\mathbf{S}}}$ in $\mathcal{I}(n_{*,\mathbf{S}})$.
6. Considering \mathcal{K}_S , Proposition 3.5 implies the existence of a new shortest path $\tilde{\gamma}_S$ which connects two nodes of $\partial\mathcal{K}_S$ with different labels.
7. We extend the path $\tilde{\gamma}_S$ as follows: If the end of γ_S has a negative label, we do nothing, and




Figure 6: The left panel shows the internal flower $\mathcal{I}(n_*)$ of n_* . We split it in succession along the green, red and blue paths. We first split n_* into $n_{*,R}$ and $n_{*,L}$ along the green path. The middle panel shows $\mathcal{I}(n_{*,R})$ and the green node is $n_{*,L}$. The shaded region is \mathcal{K}_R and the new labels are L. One end of the red path has a negative label, while the other has the label L and must therefore be connected to $n_{*,L}$. We obtain the cutting path γ_R , and after the cut, we obtain two pieces \mathcal{K}_{RR} and \mathcal{K}_{RL} . In the third panel we show the hemisphere of $n_{*,RR}$. The blue path has labels L and RL at its extremities, which must therefore be connected to $n_{*,L}$ and $n_{*,RL}$. This defines the cutting path γ_{RR} . Note that it always suffices to add at most 2 dashed segments.

if the label is some sequence \mathbf{S}' we connect the end to $n_{*,\mathbf{S}'}$ by one edge. Doing this for both ends we obtain a path $\gamma_{\mathbf{S}}$.

8. Perform a split-a-node-along-a-path on $\gamma_{\mathbf{S}}$ and continue with step 4 for the pieces $\mathcal{K}_{\mathbf{S}L}$ and $\mathcal{K}_{\mathbf{S}R}$.

Remark 4.7. The boundary of a hemisphere $\mathcal{I}_{\mathbf{S}} = \mathcal{I}(n_{*,\mathbf{S}})$ is composed of 2 types of nodes:

1. Nodes with a negative label which are part of the original boundary \mathcal{E} .
2. The children $n_{*,\mathbf{S}'}$ of the original node n_* , where \mathbf{S}' is a sequence of R's and L's.

The boundary of a piece $\mathcal{K}_{\mathbf{S}}$, which is a sub-triangulation of $\mathcal{I}_{\mathbf{S}}$, is also composed of 2 types of nodes:

1. Nodes with a negative label which are part of the original boundary \mathcal{E} , and therefore they are part of the boundary of the hemisphere $\partial\mathcal{I}_{\mathbf{S}}$ as well.
2. The other nodes whose label is some sequence \mathbf{S}' . These nodes satisfy the following two conditions:
 - (a) All nodes of $\partial\mathcal{K}_{\mathbf{S}}$ with the same label form a connected arc of $\partial\mathcal{K}_{\mathbf{S}}$.
 - (b) If a node $y \in \partial\mathcal{K}_{\mathbf{S}}$ has the label \mathbf{S}' , then y is an internal node of the hemisphere $\mathcal{I}_{\mathbf{S}}$ seen as a 2d triangulation. Furthermore, $n_{*,\mathbf{S}'} \in \partial\mathcal{I}_{\mathbf{S}}$ and $(y, n_{*,\mathbf{S}'})$ is an internal edge of the triangulation $\mathcal{I}_{\mathbf{S}}$.

Theorem 4.8. The algorithm decomposes the triangulation \mathcal{I} , and the sub-pieces \mathcal{K}_L , \mathcal{K}_R by sequences of paths γ_{s_1, \dots, s_k} until all the pieces $\mathcal{K}_{s_1, s_2, \dots, s_k}$ with $s_i \in \{R, L\}$ are reduced to simple triangles. Furthermore,

every edge in $\mathcal{I} \setminus \partial\mathcal{I}$ is in at most one path.

Proof. We need to check that the different steps of the algorithm can be performed. The steps 1–3 follow from the definition of split-a-node-along-a-path. Steps 4 and 5 need no verification. Step 6 relies on Proposition 3.5, which implies the existence of a (shortest) path $\tilde{\gamma}_S$, cutting the admissible piece \mathcal{K}_S into two admissible pieces \mathcal{K}_{SL} and \mathcal{K}_{SR} .

In step 7, we need to make sure that the path γ_S connects two *different* nodes of $\mathcal{E}(n_{*,S})$ which is also $\partial\mathcal{I}(n_{*,S})$, to be distinguished from $\partial\mathcal{K}(n_{*,S})$. The whole construction of labels has been done with this aim in mind. Note that if a node u has a negative label, we do nothing because any node u with a negative label is part of the original boundary $\mathcal{E}(n_*)$, implying that if $u \in \mathcal{I}(n_{*,S})$, then $u \in \mathcal{E}(n_{*,S})$ for any child $n_{*,S'}$ of n_* . On the other hand, if the label is the sequence S' , then by construction (step 5), u is connected to $n_{*,S'}$ with one edge. Since the labels are different by construction, the path γ is a splitting path, and therefore a cut along it is possible. In step 8, we need to verify that the cut can indeed be done, and that the algorithm can be applied to the children of the \mathcal{K} which was just cut. But this is the content of Proposition 3.5, which shows that the cut can be done in such a way that the children are admissible in the sense of Definition 3.3.

Since new paths are always constructed in the interior of \mathcal{K} , and the \mathcal{K} 's are cut along them, it is obvious that no edge is covered by more than one path. □

4.3.3 Reducing C1-nodes to C0-nodes

Let T be a triangulation of a ball. Consider an external node n_* of T and let $\mathcal{I} = \mathcal{I}(n_*)$ be its internal hemisphere. Furthermore, assume that all nodes of \mathcal{I} are either external (with regard to T) or internal of type C0 or C1 but not C2. We will describe an algorithm which promotes all the internal nodes of type C1 of \mathcal{I} to internal nodes of type C0. The approach is somewhat different from that of the previous section. Indeed promoting an internal node x of type C2 to an internal node of type C1 is done by splitting some external node n_* along a path going through x . However, let $x \in \mathcal{I}(n_*)$ be an internal node of type C1 and let (x, y, n_*) be an internal face which defines x as C1; by hypothesis, $y \in \partial\mathcal{I}$. Promoting x to an internal node of type C0 is done by splitting n_* along a path which contains the edge (y, x) .

For every internal node x of type C1 in $\mathcal{I}(n_*)$ we choose one of the $y \in \partial\mathcal{I}$ for which (x, y, n_*) is an internal face and call it $y(x)$. We define

$$\mathcal{Y} = \{(x, y(x)) \mid x \text{ is C1} \}.$$

We will eliminate elements in the list \mathcal{Y} by iterating an algorithm similar to the one in the previous section, until none are left. A binary tree of left and right pieces will be formed in the process (see Fig. 5).

At the first step of this algorithm, this tree only contains one element, namely the hemisphere \mathcal{I} . We will form a tree of \mathcal{K} 's as before, starting at $\mathcal{K} = \mathcal{I}$.

The algorithm starts with steps 1 and 2 below, and then repeats the other steps until it stops.

1. Pick an edge $(x, y) = (x, y(x)) \in \mathcal{Y}$.
2. By hypothesis, $y \in \partial\mathcal{I}(n_*)$. By Lemma 3.2 there is a second, disjoint, simple path connecting x to a node $z \in \partial\mathcal{I}(n_*)$, $z \neq y$. This defines a splitting path γ connecting 2 distinct nodes y

and z of $\partial\mathcal{I}(n_*)$. Similarly to the previous section, we split n_* along γ into $n_{*,R}$ and $n_{*,L}$. We add the 2 new pieces $\mathcal{K}(n_{*,R})$ and $\mathcal{K}(n_{*,L})$ as two leaves of \mathcal{K} in the tree. We remove the edge (x, y) from the list \mathcal{Y} . Note that the path γ might promote a second internal node x' of type C1 into a node of type C0, if the edge (x', z) is in the list \mathcal{Y} and in the path γ . In that case, both edges (x, y) and (x', z) are removed from \mathcal{Y} .

3. If the list \mathcal{Y} is empty, we are done.
4. Pick an edge $(x, y) \in \mathcal{Y}$.
5. Find the piece $\mathcal{K}(n_{*,s_1,\dots,s_k})$, where $s_i \in \{L, R\}$, among the leaves of the binary tree which contains the edge (x, y) . We use the abbreviations $\mathbf{s} = \{s_1, \dots, s_k\}$ and $n_{*,\mathbf{s}}$. The edge (x, y) belongs to exactly one piece³.
6. Observe that the node y is in $\partial\mathcal{I}(n_{*,\mathbf{s}}) \cap \partial\mathcal{I}(n_*)$ ⁴.
 - If x is in the interior of $\mathcal{K}(n_{*,\mathbf{s}})$, the edge (x, y) gives us the first simple path connecting x to $\partial\mathcal{I}(n_{*,\mathbf{s}})$ and by Lemma 3.2 there is a second independent path connecting x to a node $z \in \partial\mathcal{K}(n_{*,\mathbf{s}})$, $z \neq y$.
 If z is also in $\partial\mathcal{I}(n_{*,\mathbf{s}})$ we have found a $\gamma_{\mathbf{s}}$ along which we can cut. Note that in this case, the path $\gamma_{\mathbf{s}}$ might promote a second node x' of type C1; this happens if $z \in \partial\mathcal{I}(n_*)$ and (x', z) is an edge of $\gamma_{\mathbf{s}}$.
 If $z \notin \partial\mathcal{I}(n_{*,\mathbf{s}})$, the path $\gamma_{\mathbf{s}}$ is obtained by adding the edge which connects z to the tip of the cone⁵.
 - If x is not in the interior of $\mathcal{K}(n_{*,\mathbf{s}})$, $\gamma_{\mathbf{s}}$ is found by connecting x to a tip of one of the cones attached to $\mathcal{K}(n_{*,\mathbf{s}})$ ⁶ (see Footnote 5).
7. We split $n_{*,\mathbf{s}}$ along the path $\gamma_{\mathbf{s}}$ and add the 2 new pieces $\mathcal{K}(n_{*,\mathbf{s}R})$ and $\mathcal{K}(n_{*,\mathbf{s}L})$ to the tree as leaves of $\mathcal{K}(n_{*,\mathbf{s}})$. Note that $\mathcal{K}(n_{*,\mathbf{s}})$ is no longer a leaf of the tree and will never be encountered in the remaining steps of the algorithm. Finally, we remove the edge (x, y) (and eventually (x', z) if x' is also promoted by $\gamma_{\mathbf{s}}$) from the list \mathcal{Y} .
8. We continue with step 3.

The algorithm stops when all internal nodes of type C1 of $\mathcal{I}(n_*)$ have been promoted to C0. Since each branch of the tree is used at most once and since we never cut along the boundary of any $\mathcal{K}_{\mathbf{s}}$ we have shown:

Theorem 4.9. *The algorithm decomposes the triangulation \mathcal{I} along a sequence of simple paths; it promotes all of the internal nodes of \mathcal{I} of type C1 into nodes of type C0. Furthermore, every edge in $\mathcal{I} \setminus \partial\mathcal{I}$ is in at most one path.*

4.3.4 Change of the f-vector after the entire recursion

In this section, we compute the total change in the f-vector resulting from the elimination of all internal nodes.

³Note that the only edges which are common to more than one piece are the edges of the paths along which we already cut. Since (x, y) is still in the list \mathcal{Y} , it cannot be such an edge.

⁴By hypothesis, $y \in \partial\mathcal{I}(n_*)$ and therefore also $y \in \partial\mathcal{I}(n_{*,\mathbf{s}}) \cap \partial\mathcal{I}(n_*)$.

⁵The distance between $\partial\mathcal{I}(n_{*,\mathbf{s}})$ and any node in $\partial\mathcal{K}(n_{*,\mathbf{s}})$ is at most 1, see Fig. 6. The node z belongs to a path $\gamma_{\mathbf{s}'}$ along which we already cut. This implies that z is connected to $n_{*,\mathbf{s}'L}$ or $n_{*,\mathbf{s}'R}$, called the tip of the cone associated with z .

⁶Note that the node y is *not* on a tip of a cone but is on the original boundary $\partial\mathcal{I}(n_*)$.

Note that we will compute the total increase in the number of internal edges instead of tetrahedra. The two numbers are related by (2.1).

Definition 4.10. *We need 3 counters at each depth d of the original triangulation:*

- a_d is the number of internal edges (x, y) with $D_x = d$ and $D_y = d + 1$.
- b_d is the number of internal edges (x, y) with $D_x = d$ and $D_y = d$.
- c_d is the number of internal faces (x, y, z) with $D_x = d$ and $D_y = d + 1$. (This implies $D_z = d$ or $d + 1$.)

As every node is connected to nodes of the same depth or to depths differing by at most 1, the following obvious relations hold:

$$\begin{aligned} \sum_d (a_d + b_d) &= e, \\ \sum_d c_d &\leq f_i, \end{aligned} \tag{4.1}$$

where e is the number of internal edges, and f_i is the number of internal faces.

Let Δ_d denote the increase of the number of internal edges obtained when performing the steps $C2 \rightarrow C1 \rightarrow C0 \rightarrow \text{external}$ at level d .

Proposition 4.11. *There is a constant C'_Δ such that*

$$\begin{aligned} \Delta_d &\leq C'_\Delta (a_d + b_d + a_{d-1} + c_{d-1}), \text{ for } d > 0, \\ \Delta_0 &\leq C'_\Delta (a_d + b_d + n_s), \text{ for } d = 0. \end{aligned} \tag{4.2}$$

Corollary 4.12. *Eliminating all internal nodes of a triangulation T with f -vector $\langle t, f_s, n_i \rangle$ leads to a total increase Δ of internal edges which is bounded by*

$$\Delta \leq C_\Delta (t + n_i).$$

Proof of the corollary. From (2.1) we deduce $f_i = 2t - f_s/2$ and $e = t + n_i - f_s/2 + 1$. Also, $n_s = f_s/2 + 2$. Using (4.1) and the proposition, we get

$$\Delta = \sum_{d \geq 0} \Delta_d \leq C'_\Delta (2e + f_i + n_s),$$

from which the assertion follows (the coefficient of f_s is negative and the additive constants can be bounded since $1 \leq t$). \square

Proof of Proposition 4.11. The key to the bound (4.2) is the observation that the transformations $C2 \rightarrow C1 \dots$ are *local* in the depth d one considers. Indeed, as is visible from the definition of these transformations, working at level d only affects a_d, b_d, c_d and a_{d-1}, c_{d-1} .

More precisely, when starting to work at level d , we need the value of \hat{a}_{d-1} , which is the number of internal edges (connecting depth $d - 1$ to d) obtained when level $d - 1$ has been completed.

As we work on level d , these values continue to change. After the sweep $C2 \rightarrow C1$ at level d we obtain a'_d, \hat{a}'_{d-1} , and similarly for the other variables. After the sweep $C1 \rightarrow C0$ we obtain a''_d and other variables. The sweep of removing the tetrahedra after $C0$ decreases all the counters, so we do not introduce new notation.

The main bound is

Lemma 4.13. *One has, after completing level $d - 1$:*

$$\hat{a}_{d-1} \leq a_{d-1} + 16c_{d-1} . \quad (4.3)$$

Postponing the proof, we recall the following facts:

- Cutting along a path γ adds $|\gamma| - 1$ internal edges to the triangulation (see Lemma 4.6).
- Each edge of each $\mathcal{I}(n_*)$ is used in at most 1 path $\tilde{\gamma}_s$ (see Theorem 4.8 and 4.9).
- The extension of the path $\tilde{\gamma}_s$ to a splitting path γ_s adds at most 2 to its length (see step 7 for the case $C2 \rightarrow C1$, and step 6 for the case $C1 \rightarrow C0$). We will use this observation by saying that $|\gamma_s| - 1 \leq 2|\tilde{\gamma}_s|$.

Using these facts and Lemma 3.1, the increase Δ'_d of the number of internal edges due to the sweep $C2 \rightarrow C1$ is bounded by

$$\begin{aligned} \Delta'_d &\leq 2 \sum_{n_* \in \mathcal{L}_d} \sum_{\mathbf{s}} |\tilde{\gamma}_s| \leq 2 \sum_{n_* \in \mathcal{L}_d} \#(\text{edges in } \mathcal{I}(n_*)) \\ &\leq 6a_d + 12b_d + 6\hat{a}_{d-1} + 2 \sum_{n_* \in \mathcal{L}_d} (|\mathcal{E}(n_*)| - 3) , \end{aligned} \quad (4.4)$$

Here, we over-count the number of added internal edges. However, one should keep in mind that if we follow the algorithms of Sections 4.3.2 and 4.3.3, then the Theorems 4.8 and 4.9 are valid and every new internal edge is accounted for. As a consequence, the Relation (4.4) is an upper bound on the number of internal edges due to the sweep $C2 \rightarrow C1$.

The effect of the sweep $C2 \rightarrow C1$ at level d is summarized by

Lemma 4.14. *One has*

$$a'_d + b'_d + a'_{d-1} \leq a_d + b_d + \hat{a}_{d-1} + \Delta'_d , \quad (4.5a)$$

$$a'_d \leq a_d + 2c_d , \quad (4.5b)$$

$$c'_d \leq 7c_d . \quad (4.5c)$$

Postponing the proof, we proceed to the sweep $C1 \rightarrow C0$. In the same manner, the increase Δ''_d of internal edges for the sweep $C1 \rightarrow C0$ at level d is

$$\begin{aligned} \Delta''_d &\leq 2 \sum_{n_* \in \mathcal{M}_d} \sum_{\mathbf{s}} |\tilde{\gamma}_s| \leq 2 \sum_{n_* \in \mathcal{M}_d} \#(\text{edges in } \mathcal{I}(n_*)) \\ &\leq 6a'_d + 12b'_d + 6a'_{d-1} + 2 \sum_{n_* \in \mathcal{M}_d} (|\mathcal{E}(n_*)| - 3) . \end{aligned} \quad (4.6)$$

To complete the proof of Proposition 4.11 we note that the external degree of n_* is always 3 for those nodes which have been promoted to the surface by removing a tetrahedron. (Those which were at the surface at level $d = 0$ can of course have higher degree.) Using (4.4) and Lemma 4.13 we get

$$\begin{aligned} \Delta'_d &\leq 6a_d + 12b_d + 6\hat{a}_{d-1} + 4e_s \cdot \delta_{d=0} \\ &\leq 6a_d + 12b_d + 6a_{d-1} + 96c_{d-1} + 12n_s \cdot \delta_{d=0} \\ &\leq 96(a_d + b_d + a_{d-1} + c_{d-1} + n_s \cdot \delta_{d=0}) . \end{aligned} \quad (4.7)$$

In (4.6) the external degree of a node $n_* \in \mathcal{M}_d$ can be larger than 3. However, if we split a node n_* into $n_{*,R}$ and $n_{*,L}$, then the external degrees satisfy

$$|\mathcal{E}(n_*)| = |\mathcal{E}(n_{*,L})| + |\mathcal{E}(n_{*,R})| - 4. \quad (4.8)$$

Therefore, we can bound

$$\sum_{n_* \in \mathcal{M}_d} (|\mathcal{E}(n_*)| - 3) \leq 4 \cdot \#(\text{splits in C2} \rightarrow \text{C1}).$$

Since each split adds at least one internal edge, we deduce that

$$\sum_{n_* \in \mathcal{M}_d} (|\mathcal{E}(n_*)| - 3) \leq 4\Delta'_d.$$

Combining this with (4.5a) and Lemma 4.13, we get

$$\begin{aligned} \Delta''_d &\leq 6a'_d + 12b'_d + 6a'_{d-1} + 2 \sum_{n_* \in \mathcal{M}_d} (|\mathcal{E}(n_*)| - 3) \\ &\leq 6a_d + 12b_d + 6a_{d-1} + 96c_{d-1} + 20\Delta'_d. \end{aligned}$$

Replacing Δ'_d with (4.7) yields the result we seek. The last step C0 \rightarrow external adds no internal edges; in fact it reduces their number. This finishes the proof of Proposition 4.11. \square

Proof of Lemma 4.14. In the sweep C2 \rightarrow C1 at level d , we split all (or some) nodes $\{n_*\} \subset \mathcal{L}_d$ into $\{n_{*,s}\} \subset \mathcal{M}_d$. As a consequence, all Δ'_d added internal edges have an end $n_{*,s}$ which is the child of some node $n_* \in \mathcal{L}_d$ at depth d in the initial triangulation. The number of internal edges having a corner at depth d is given by $a_d + a_{d-1} + b_d$. This proves the relation (4.5a).

To prove (4.5b), we need to bound the added number of internal edges $(n_{*,s}, y)$ such that $n_* \in \mathcal{L}_d$ and $y \in \mathcal{I}(n_*) \cap \mathcal{L}_{d+1}$ was at depth $D_y = d+1$ in the original triangulation (and therefore is at depth 1 in the current step). By construction, this number is bounded by the number of paths γ_s which go through such a node y in the 2d triangulation $\mathcal{I}(n_*)$. Furthermore, by Theorem 4.8, each edge of $\mathcal{I}(n_*)$ is used in at most one path γ_s . We deduce that, for two such nodes n_* and y , the number of added internal edges of type $(n_{*,s}, y)$ is bounded by the degree of the edge (n_*, y) in the original triangulation. Summing the degrees of all edges (n_*, y) such that $n_* \in \mathcal{L}_d$ and $y \in \mathcal{I}(n_*) \cap \mathcal{L}_{d+1}$ is bounded by $2c_d$.

Finally, in order to prove (4.5c), we need to bound the added number of internal faces $(n_{*,s}, y, z)$ in the step C2 \rightarrow C1 at level d when $n_{*,s}$ is obtained from splitting some $n_* \in \mathcal{L}_d$ and $y \in \mathcal{L}_{d+1} \cap \mathcal{I}(n_*)$. But each added internal face $(n_{*,s}, y, z)$ requires the addition of the internal edge $(n_{*,s}, y)$. Furthermore, by definition of the move split-a-node-along-a-path, each new internal edge is added along with three internal faces. We deduce that $c'_d - c_d \leq 3(a'_d - a_d) \leq 6c_d$. \square

Proof of Lemma 4.13. The proof follows by induction on the level d . At level $d - 1 = 0$, there is nothing to prove (since no splits have been done). When we are at level $d - 1 > 0$ we can use Proposition 4.11 (at level $d - 1$). Following the same reasoning used in the proof of (4.5b), we can write

$$a''_{d-1} \leq a'_{d-1} + 2c'_{d-1}.$$

Replacing (4.5b) and (4.5c), we get

$$a''_{d-1} \leq a_{d-1} + 16c_{d-1} .$$

Finally, $\hat{a}_{d-1} \leq a''_{d-1}$ since the third step $C0 \rightarrow \text{external}$ does not add internal edges (this third step actually removes $3|\mathcal{L}_{d-1}|$ internal edges).

□

4.4 Reducing a triangulation with no internal nodes into a set of nuclei

Let T be any triangulation. In the previous section, we described an algorithm which transforms T into a new triangulation T' with no internal nodes. We now systematically apply the moves cut-a-3-face and open-a-2-face on every internal face of T' with less than 2 internal edges. We end up with a collection of triangulations $\{N_i\}$ satisfying the following properties:

- All nodes of any such N_i are external.
- All internal faces of any such N_i have at least 2 internal edges.

Any triangulation satisfying these two conditions is called a *nucleus*.

5 Part II: Bounding the number of triangulations

We showed that any triangulation can be reduced into a collection of nuclei using four moves. For the moment, we proceed without using the move cut-a-3-face. This implies that any triangulation can be transformed into a “tree of nuclei” (the formal definition of a tree of nuclei will be given later on) using the three remaining moves. Equivalently, this shows that any triangulation can be constructed from a tree of nuclei, using the inverse of these three moves. Bounding the number of trees of nuclei, and then bounding the number of ways one can perform the inverse moves on such a tree yields a bound on the total number of triangulations.

5.1 Rooted triangulations

We define what we mean by a rooted triangulation T and we show that one can label all external nodes of T . In the sequel, we use a particular labeling described below.

Definition 5.1. A rooted triangulation (T, F) of the 3-ball is a triangulation T with one labeled external face F . This labeled face is called the root. The three nodes of the root are always labeled 0, 1, and 2.

Remark 5.2. We will only consider rooted triangulations. This means for instance that talking about the Christmas tree $T_m, m > 1$ makes no sense, since there is more than one such rooted triangulation. The exceptions are of course symmetric triangulations T such as the tetrahedron.

Proposition 5.3. Consider the boundary of a rooted triangulation (T, F) . The root is labeled as $(0, 1, 2)$. One can define a way of labeling all external nodes of T .




Figure 7: The Christmas tree with $m = 3$ internal nodes. This triangulation can be rooted in more than one way.

Proof. The proof is just the construction of this labeling. Any labeled edge can be seen as an element $(a, b) \in \mathbb{Z}_+^2$ with $a < b$.⁷ We consider the lexical order on \mathbb{Z}_+^2 . We start with the node 0. Its external flower is a 1d triangulation of the circle S^1 and it contains the edge $(1, 2)$ by definition. This edge determines the direction in which we label all unlabeled nodes of the external flower of node 0.

Next, we consider the external flower of node 1 and we look for the smallest labeled edge in the sense of the above ordering. In this case, this edge is $(0, 2)$. This edge fixes the direction in which we label all the yet unlabeled nodes of the external flower of node number 1. Notice that all unlabeled nodes which are assigned a label are part of a face along with 2 already labeled nodes. This implies that the external flower of any labeled node contains a smallest labeled edge and as such can be directed.

We continue with all the nodes in their natural order until all external nodes of T are labeled. \square

5.2 Trees of nuclei

Since we work with rooted triangulations, from now on, we will only use rooted nuclei, namely:

Definition 5.4. A nucleus is a **rooted** triangulation with no internal nodes such that every internal face has at most one external edge.

⁷We use the notation $\mathbb{Z}_+ = \{0, 1, 2, \dots\}$.

5.2.1 Rooted trees of nuclei and planar rooted trees

Let \mathcal{N} be the set of all nuclei and $\mathcal{N}_{t,f}$ be the subset of all nuclei with t tetrahedra and f external faces.

Definition 5.5. *A rooted triangulation T is called a rooted tree of nuclei if all nodes of T are external and all internal faces of T have 0, 2, or 3 internal edges. (In other words, no internal face has 2 external edges.)*

In other words, a rooted tree of nuclei is simply a rooted triangulation which is obtained by gluing sequentially nuclei along pairs of their external faces. This is done in such a way that each nucleus is glued to an external face (a, b, c) of its parent through its root; 0 is identified with a , 1 with b and 2 with c . Once the tree is built, the external nodes are renumbered in the sense of Proposition 5.3.

Since all external faces of a rooted triangulation are ordered, this defines a bijection between rooted trees of nuclei (T, F) and rooted planar trees with colored vertices in the following manner:

- Each nucleus of the triangulation (T, F) is represented by a colored vertex.
- The root-vertex of the planar tree represents the nucleus with the root F , *i.e.*, with the face $(0, 1, 2)$.
- Each internal face of the triangulation with three external edges is shared by two nuclei and hence it is represented in the tree by an edge linking the corresponding two colored vertices.
- Since the internal faces with three external edges are ordered, this induces an order of the links of the planar tree, say from left to right.

5.2.2 Hypothesis on the number of rooted nuclei

We next show how the question of Gromov can be reformulated. We show that if there are not “too many” different types of nuclei, then there is indeed an exponential bound on the number of triangulations, when expressed in terms of the number of tetrahedra.

Hypothesis 5.6. *There is a finite constant $K_1 > 1$ such that the number $\varrho(t, f_s)$ of face-rooted nuclei with f -vector $\langle t, f_s, 0 \rangle$ is bounded by K_1^t .*

In order to alleviate the notation, from now on, we will denote f_s by f .

Lemma 5.7. *For any nucleus $N \in \mathcal{N}_{t,f}$ one has $f \leq t + 3$.*

Proof. If N is a tetrahedron, the assertion is obvious. If N is non-trivial each tetrahedron of N can have at most 1 external face, since otherwise it would have an internal face with more than one external edge. \square

5.2.3 The number of rooted trees of nuclei

We use the classical method for counting planar ordered trees, generalized to the case of a multitude of different nodes, which are the face-rooted nuclei.

Definition 5.8. *Let $A_{v,t,f}$ be the number of rooted trees of nuclei with $v > 0$ nuclei, t tetrahedra and f external faces. We define $A_{0,t,f} = \delta_{t,0} \delta_{f,0}$.*

Our main bound is:

Proposition 5.9. *Under the Hypothesis 5.6 there is a K_2 , with $2 < K_2 < \infty$ such that for all t, f , one has*

$$\sum_v A_{v,t,f} \leq K_2^t.$$

Proof. Consider a tree of nuclei, and let N be the nucleus containing the root F and assume that $N \in \mathcal{N}_{t_0, f_0}$. Removing N from the tree leads to $f_0 - 1$ rooted trees of nuclei, some of which may be empty. We let v_i, t_i , and f_i denote the counters for the branch i . Note that if a branch i has 0 nuclei, *i.e.*, if $v_i = 0$, then, obviously, $t_i = f_i = 0$. Thus, we get the relations:

$$\sum_{i=1}^{\ell} v_i = v - 1, \quad \sum_{i=1}^{\ell} t_i = t - t_0, \quad \sum_{i=1}^{f_0-1} \delta_{f_i > 0} (f_i - 1) + \delta_{f_i = 0} = f - 1. \quad (5.1)$$

In the sequel, we denote by \sum'_{v,t,f,t_0,f_0} the sum over the set

$$\{v_i, t_i, f_i \mid i = 1, \dots, f_0 - 1, v_i \geq 0, t_i \geq 0, f_i \geq 0 \text{ and satisfying (5.1)}\}.$$

This observation allows us to write a recursive relation

$$A_{v,t,f} = \delta_{v,0} \delta_{t,0} \delta_{f,0} + \sum_{t_0 > 0, f_0 \geq 4} \varrho(t_0, f_0) \sum'_{v,t,f,t_0,f_0} \prod_{i=1}^{f_0-1} A_{v_i, t_i, f_i}. \quad (5.2)$$

Fix $M \in \mathbb{Z}_+$, and assume that v, t, f satisfy $3v + 3t + f \leq M$. By (5.1), we deduce

$$3v_i + 3t_i + f_i \leq 3v - 3 + 3t - 3t_0 + f \leq M - 1.$$

We define

$$A_M(s) = \sum_{3v+3t+f \leq M} A_{v,t,f} s^{3v+3t+f}.$$

Clearly, $A_0(s) = 1$ for all s , $A_M(0) = 1$ for all $M \geq 0$, and for a fixed s , $A_M(s)$ is an increasing sequence in M .

Multiplying (5.2) by $s^{3v+3t+f}$ and summing, we get, using (5.1):

$$\begin{aligned} A_M(s) &= 1 + \sum_{3v+3t+f \leq M} \sum_{t_0=1}^t \sum_{f_0=4}^f \varrho(t_0, f_0) s^{3+3t_0+1-\sum_{i=1}^{f_0-1} (\delta_{f_i > 0} - \delta_{f_i = 0})} \\ &\quad \times \sum'_{v,t,f,t_0,f_0} \prod_{i=1}^{\ell} A_{v_i, t_i, f_i} s^{3v_i+3t_i+f_i}. \end{aligned} \quad (5.3)$$

Using Lemma 5.7, we have

$$\begin{aligned} 3 + 3t_0 + 1 - \sum_{i=1}^{f_0-1} (\delta_{f_i > 0} - \delta_{f_i = 0}) &\geq 3 + 3t_0 + 1 - (f_0 - 1) \cdot 1 + 0 \\ &\geq 5 + 3t_0 - f_0 = 2(t_0 + 3 - f_0) + t_0 + f_0 - 1 \end{aligned}$$

$$\geq t_0 + f_0 - 1 .$$

Restricting to $0 \leq s \leq 1$, this implies

$$s^{3+3t_0+1-\sum_{i=1}^{f_0-1}(\delta_{f_i>0}-\delta_{f_i=0})} \leq s^{t_0+f_0-1} . \quad (5.4)$$

Using now the Hypothesis 5.6, *i.e.*, $\varrho(t, f) \leq K_1^t$, we get from (5.3) and (5.4):

$$A_M(s) - A_M(0) \leq \sum_{t_0=0}^M (sK_1)^{t_0} \sum_{f_0-1=0}^M \prod_{i=1}^{f_0-1} sA_{M-1}(s) \leq \frac{1 - (sK_1)^{M+1}}{1 - sK_1} \frac{1 - (sA_{M-1}(s))^{M+1}}{1 - (sA_{M-1}(s))} .$$

Restricting s further to $s \leq 1/(2K_1)$ this leads to

$$A_M(s) - A_M(0) \leq 2 \frac{1 - (sA_{M-1}(s))^{M+1}}{1 - (sA_{M-1}(s))} .$$

Fix $s^* = \min(0.1, 1/(2K_1))$ and consider the map $F : x \mapsto 1 + 2/(1 - s^* \cdot x)$. One easily checks that F maps the interval $[1, 5]$ to itself. Furthermore, we have $s^* \cdot x \leq 1$ for $x \in [1, 5]$. Starting with $x = A_0(s^*) = 1$ we conclude that for all M one has $A_M(s^*) \leq 5$. This implies that the monotone sequence $A_M(s^*)$ converges as $M \rightarrow \infty$ and thus

$$A_{v,t,f} \leq 5 \cdot (s^*)^{-3v-3t-f} .$$

Summing over v and using $v \leq t$ and $f \leq 4t$ we complete the proof. \square

5.3 Bound on triangulations

Having discussed the number of trees, we now study the number of ways these trees can be made into triangulations by identifying faces and nodes. This process is patterned after the work of [7] and [4].

Our bounds are based on using the inverses of the moves open-a-2-face, remove-1-tetra, and split-a-node-along-a-path. Since we are only interested in the bound, we will allow for inverse moves which do not necessarily lead to 3-balls.

Remark 5.10. *While we over-count the number of triangulations, by allowing for moves which may not lead to 3-balls, we can in fact formulate precise conditions which guarantee that after each move, a 3-ball is obtained. These conditions are spelled out in Lemmas 5.11 and 5.15. This actually allows for efficient programming of the inverse operations.*

5.3.1 Bounding the number of rooted triangulations with no internal nodes

Let $\mathcal{R}_{t,f}$ be the set of all rooted trees of nuclei with t tetrahedra and f external faces and let $\mathcal{T}_{t,f,0}$ be the set of all rooted triangulations with t tetrahedra, f external faces and no internal nodes. In this section, we will define the inverse move of open-a-2-face and we will use it to count the number of rooted triangulations with no internal nodes.

The inverse operation of open-a-face, which we will simply call *identification* when there is no ambiguity, is to identify two adjacent external faces, satisfying some conditions. Indeed, identifying

any two adjacent external faces might lead to a complex which is not a triangulation. For instance, assume that (n_1, n_2, m_1) and (n_1, n_2, m_2) are two adjacent external faces such that there exists a node x adjacent to both m_1 and m_2 . After identifying the two faces, we obtain a complex with a double edge $(x, m_1) = (x, m_2)$.

Lemma 5.11. *Consider a triangulation T . Let (a, b) be an external edge and let x, y be its two opposite external nodes. Assume that the following conditions are satisfied:*

- *The nodes x and y are not connected by an edge.*
- *The only nodes m such that (m, x) and (m, y) are edges are the two nodes a and b .*

Then, one can identify the two external nodes x and y as well as the two external faces sharing (x, y) . This operation transforms a 3-ball to a 3-ball, and will be called identification (of two adjacent external faces).

Proof. The proof is left to the reader. □

Proposition 5.12. *Under Hypothesis 5.6, there is a constant K_3 such that for all t and f one has*

$$|\mathcal{T}_{t,f,0}| \leq K_3^t.$$

Proof. Let $T \in \mathcal{T}_{t,f,0}$ be any rooted triangulation with no internal nodes. Using repetitively the move open-a-2-face on T transforms it into a rooted triangulation T' with no internal nodes such that each internal face has 0, 1 or 3 external edges. In other words, T' is a rooted tree of nuclei. Equivalently, given a rooted tree of nuclei T' with t' tetrahedra and f' external faces, one can count the number of ways one can identify two adjacent external faces, *without any conditions guaranteeing ballness*. Multiplying this number by the number of rooted trees of nuclei gives us an upper bound on the number of rooted triangulations with no internal nodes.

We count the number of $T \in \mathcal{T}_{t,f,0}$ obtained by identification from a rooted tree of nuclei T' with t' tetrahedra and f' external faces. This means that we identify $D = (f' - f)/2$ pairs of adjacent external faces.

We first observe that choosing a pair of adjacent external faces is equivalent to choosing an external edge. We then note that some faces which are not adjacent in T' might become adjacent after some identifications are done. This means that we have a sequence e_1, e_2, \dots, e_ℓ with $e_i \geq 1$ and $\sum_i e_i = D$ which is defined as follows:

- e_1 is the number of external edges (or equivalently of pairs of adjacent external faces) of T' which are identified.
- e_2 is the number of pairs of faces which were not adjacent in T' but became so after the first series of e_1 identifications. However, each identification of two adjacent external faces creates exactly two new pairs of adjacent external faces, implying that $e_2 \leq 2e_1$.
- e_i is defined by analogy from the e_{i-1} identifications, implying that $e_i \leq 2e_{i-1}$.

This leads to the following bound:

$$|\mathcal{T}_{t,f,0}| \leq \sum_{f' > f} |\mathcal{R}_{t,f'}| \sum_{\ell=1}^{D \equiv (f' - f)/2} \sum_{\sum_{i=1}^{\ell} e_i = D, e_i \geq 1} \binom{3D}{e_1} \binom{2e_1}{e_2} \cdots \binom{2e_{\ell-1}}{e_\ell}.$$

Since $\binom{a}{b} \leq 2^a$, and since the number of external faces f' in any rooted tree of nuclei is bounded by four times the number of tetrahedra, we find, using Proposition 5.9 to bound $|\mathcal{R}_{t,f'}|$,

$$\begin{aligned}
|\mathcal{T}_{t,f,0}| &\leq \sum_{f' > f} |\mathcal{R}_{t,f'}| 2^{5(f'-f)/2} \sum_{\ell=1}^{D \equiv (f'-f)/2} \sum_{\sum_{i=1}^{\ell} e_i = D, e_i \geq 1} 1 \\
&\leq \sum_{f' > f} |\mathcal{R}_{t,f'}| 2^{5(f'-f)/2} \sum_{\ell=1}^{D \equiv (f'-f)/2} \binom{D-1}{\ell-1} \\
&\leq \sum_{f' > f} |\mathcal{R}_{t,f'}| 2^{3(f'-f)} \\
&\leq \sum_{f'=f+2}^{4t} K_2^t K_2^{3(f'-f)} \\
&\leq K_2^{13t} = K_3^t,
\end{aligned}$$

where $K_3 = K_2^{13}$.

The proof is complete. \square

5.3.2 Bounding the number of rooted triangulations (internal nodes included)

In this section, we define the inverse moves of remove-1-tetra and split-a-node-along-a-path and we use them to count the number of rooted triangulations.

Definition 5.13. *We define the inverse move of remove-1-tetra, which we call adding a tetrahedron: Consider a triangulation T . Let x be an external node with external degree equal to 3 and let a_1, a_2 and a_3 be its external neighbors, i.e., (x, a_i) is an external edge. Adding a tetrahedron then consists in adding the face (a_1, a_2, a_3) and the tetrahedron (x, a_1, a_2, a_3) .*

We define the inverse move of split-a-node-along-a-path.

Lemma 5.14. *Consider a triangulation T . Let (a, b) be an external edge. Assume that the following conditions are satisfied:*

- *For each node m such that (m, a) and (m, b) are edges, (m, a, b) is a face.*
- *For each edge e such that (e, a) and (e, b) are faces, (e, a, b) is a tetrahedron.*
- *There are no faces f such that (f, a) and (f, b) are both tetrahedra.*

Then, one can collapse the two nodes a and b , and the result is again a 3-ball. This move is called collapse of an external edge or simply collapse.

Proof. The proof is left to the reader. \square

The three conditions of a collapse can be reformulated in the following manner:

Lemma 5.15. *Let $e = (a, b)$ be an external edge. The edge e is collapsible if and only if*

$$\mathcal{I}(a) \cap \mathcal{I}(b) = \mathcal{I}(e),$$

where $\mathcal{I}(a)$ is the hemisphere of a and $\mathcal{I}(e)$ is the semi-circular flower of e .

Proof. By definition, an edge $e = (a, b)$ is collapsible if and only if

- For each node m such that (m, a) and (m, b) are edges, (m, a, b) is a face.
- For each edge e such that (e, a) and (e, b) are faces, (e, a, b) is a tetrahedron.
- There are no faces f such that (f, a) and (f, b) two tetrahedra.

Any graph is defined as a set of vertices and a set of edges. A 2d triangulation is a graph that can be defined as a set of nodes, a set of edges and a set of faces, and a 1d triangulation as a set of nodes and a set of edges. $\mathcal{I}(a)$ is a 2d triangulation and $\mathcal{I}(e)$ is a 1d triangulation. Let $\mathcal{V}(a)$, $\mathcal{L}(a)$ and $\mathcal{F}(a)$ be the sets of vertices, edges and faces of $\mathcal{I}(a)$ and $\mathcal{V}(e)$, $\mathcal{L}(e)$ those of $\mathcal{I}(e)$. The proposition is equivalent to the following

$$\begin{aligned}\mathcal{V}(a) \cap \mathcal{V}(b) &= \mathcal{V}(e) , \\ \mathcal{L}(a) \cap \mathcal{L}(b) &= \mathcal{L}(e) , \\ \mathcal{F}(a) \cap \mathcal{F}(b) &= \emptyset .\end{aligned}$$

The two definitions are clearly equivalent. □

In Sect. 4.3, we described an algorithm which transforms any triangulation with f-vector $\langle t, f, n \rangle$ into a triangulation with f-vector $\langle t', f', 0 \rangle$. We have the following lemma:

Lemma 5.16. *There is a constant $K_4 > 0$ such that the f-vectors $\langle t, f, n \rangle$ and $\langle t', f', 0 \rangle$ satisfy the following linear relation:*

$$t' \leq K_4 t , \quad f' \leq K_4 t , \quad (5.5)$$

Proof. Let e, e' be the number of internal edges of both triangulations. By Condition 4.12, we have $e' - e \leq C_\Delta(t + n_i)$. Using (2.1) and $f_s, n_i \leq 4t$, the result follows. □

This proves that any triangulation in $\mathcal{T}_{t,f,n}$ can be obtained from a triangulation with no internal nodes in $\mathcal{T}_{t',f',0}$ with a series of carefully chosen collapses and additions of tetrahedra, with t, f, n, t', f' satisfying (5.5).

We can now use a similar approach to that of the previous section. It is clear that choosing a triplet of external faces for the move add-1-tetrahedron is equivalent to choosing an external node x , and that choosing a couple of external nodes for collapse is equivalent to choosing an external edge.

5.4 Combining the bounds

Before we state our main result, we recall the

Hypothesis 5.6. *There is a finite constant $K_1 > 1$ such that the number $\varrho(t, f)$ of face-rooted nuclei with f-vector $\langle t, f \rangle$ is bounded by K_1^t .*

Theorem 5.17. *Under Hypothesis 5.6 one has the bound: There is a finite constant C such that the number of rooted triangulations with f-vector $\langle t, f, n \rangle$ is bounded by*

$$|\mathcal{T}_{t,f,n}| \leq C^t . \quad (5.6)$$

Proof. Consider a rooted triangulation $T \in \mathcal{T}_{t,f,n}$ with t tetrahedra, f external faces and n internal nodes. We showed that T can be obtained from a rooted triangulation $T' \in \mathcal{T}_{t',f',0}$ by a series of carefully chosen collapses and additions of tetrahedra.

Note that the algorithm of Sect. 4.3 which transforms T into T' can always be stopped when the last internal node of T is removed. This implies that, in the inverse construction we are doing now, we must start by adding tetrahedra to T' , and not by collapsing external edges. So the first step is to choose n_1 external nodes (of external degree 3) out of the $f'/2 + 2$ external nodes of T' , and to insert a tetrahedron on each of them with one tip at the node. We call this “covering the node”.

This reduces the number of external edges from $3f'/2$ to $3(f'/2 - n_1)$. Then, we choose m_1 external edges and we collapse them.

Remark 5.18. Any labeled triangulation is simply defined by the list of its tetrahedra \mathcal{L}_t . In this point of view, collapsing an external edge e is simply the operation where we remove from \mathcal{L}_t all the tetrahedra of $\mathcal{E}(e)$. Let e_1 and e_2 be two collapsible edges. The construction implies that the order in which we collapse them is irrelevant and so, the idea that we simultaneously collapse m_1 edges makes sense.

One should pay attention to the case where we collapse two edges $e_1 = (a, b_1)$ and $e_2 = (a, b_2)$ such that $(b_1, b_2) = e_3$ is an edge. In this case, all tetrahedra sharing one of the three edges are removed. Clearly, this yields the same result regardless of the order in which we collapse e_1 and e_2 .

The next step is to choose n_2 external nodes among the new possibilities which appear after performing the first series of coverings and collapses, and cover them. For each external edge e , we can associate four nodes: the two endpoints of e and the two nodes x_1, x_2 such that (x_i, e) is an external face. Assume that x is one of the n_2 chosen external nodes. The fact that x appeared after the first series implies that x is either one of the four nodes associated with one of the m_1 collapsed edges (note that these four nodes become three after the collapse), or that there is a node y among the first n_1 nodes such that (x, y) was an external edge (before covering y with a tetrahedron). But each such y has exactly 3 external neighbors. This implies that $n_2 \leq 3m_1 + 3n_1$ and the number of ways to choose these nodes is bounded by

$$\binom{3(m_1 + n_1)}{n_2}.$$

Continuing in this way, we choose m_2 external edges and we collapse them. Let e be such an edge. Again, e was not among the first m_1 edges. This implies that there must be a node x of the series of n_2 covered external nodes such that (e, x) formed an external face before covering x with a tetrahedron. But for each such x there are exactly three external edges satisfying this condition. We deduce that $m_2 \leq 3n_2$.

We continue adding tetrahedra and collapsing edges. This leads to two sequences $n_i, m_i, i = 1, \dots, \ell$, with $\ell \leq n$, satisfying:

$$\begin{aligned} 1 \leq n_i, \quad 0 \leq m_i \leq 3n_i, \quad \sum_{i=1}^{\ell} n_i &= n, \\ 1 \leq n_i \leq 3n_{i-1} + 3m_{i-1}, \quad i > 1, \\ \sum_{i=1}^{\ell} 2n_i + 2m_i + f &= f'. \end{aligned} \tag{5.7}$$

Note that some, or all, of the m_i 's might be equal to zero. Using (5.7) we get a bound

$$\begin{aligned}
|\mathcal{T}_{t,f,n}| &\leq \sum_{t',f'} |\mathcal{T}_{t',f',0}| \sum_{\ell=1}^n \sum_{\sum_{i=1}^{\ell} n_i = n, n_i \geq 1} \sum_{\sum_{i=1}^{\ell} m_i = (f'-f)/2 - n, m_i \geq 0} \\
&\quad \times \binom{f'/2 + 2}{n_1} \binom{3(n_1 + m_1)}{n_2} \cdots \binom{3(n_{\ell-1} + m_{\ell-1})}{n_{\ell}} \\
&\quad \times \binom{3f'/2}{m_1} \binom{3n_1}{m_2} \cdots \binom{3n_{\ell-1}}{m_{\ell}},
\end{aligned}$$

where the sum over t', f' is restricted by (5.5). Bounding each binomial by a power of 2 and using Proposition 5.12, (5.7) and (5.5), we get, as in the proof of Proposition 5.12,

$$|\mathcal{T}_{t,f,n}| \leq \sum_{t', f' \leq K_4 t} K_3^{t'} \leq C^t.$$

This shows (5.6) and completes the proof. □

Acknowledgements

We profited from useful discussions with G. Ziegler. This work was partially supported by the Fonds National Suisse and by an ERC FP7 “Ideas” Advanced Grant.

References

- [1] E. Aharonov, E. Bouchbinder, H. G. E. Hentschel, V. Ilyin, N. Makedonska, I. Procaccia, and N. Schupper. Direct identification of the glass transition: Growing length scale and the onset of plasticity. *Europhysics Letters* **77** (2007), 56002.
- [2] J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry*. Cambridge Monographs on Mathematical Physics (Cambridge: Cambridge University Press, 1997).
- [3] T. Aste and D. Sherrington. Glass transition in self-organizing cellular patterns. *J. Phys. A* **32** (1999), 7049–7056.
- [4] B. Benedetti and G. M. Ziegler. On locally constructible spheres and balls. *Acta Math.* **206** (2011), 205–243.
- [5] R. H. Bing. Some aspects of the topology of 3-manifolds related to the Poincaré conjecture. In: *Lectures on modern mathematics, Vol. II* (New York: Wiley, 1964), pp. 93–128.
- [6] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics* (Heidelberg: Springer, 2010), fourth edition.
- [7] B. Durhuus and T. Jónsson. Remarks on the entropy of 3-manifolds. *Nuclear Phys. B* **445** (1995), 182–192.
- [8] J.-P. Eckmann. A topological glass. *J. Stat. Phys.* **129** (2007), 289–309.
- [9] J.-P. Eckmann and M. Younan. Decay of correlations in a topological glass. *Philosophical Magazine* **92** (2012), 98–119.
- [10] R. Furch. Zur Grundlegung der kombinatorischen Topologie. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* **3** (1924), 69–88.
- [11] R. E. Goodrick. Non-simplicially collapsible triangulations of I^n . *Proc. Cambridge Philos. Soc.* **64** (1968), 31–36.
- [12] M. Gromov. Spaces and questions. *Geom. Funct. Anal.* (2000), 118–161 GAFA 2000 (Tel Aviv, 1999).
- [13] M. Hachimori. Simplicial complex library. <http://infoshako.sk.tsukuba.ac.jp/~HACHI/math/library>
- [14] M. O. Magnasco. Two-dimensional bubble rafts. *Philosophical magazine B* **65** (1992), 895–920.
- [15] J. Pfeifle and G. M. Ziegler. Many triangulated 3-spheres. *Math. Ann.* **330** (2004), 829–837.
- [16] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.* **14** (1962), 21–38.